arXiv:1609.04782v1 [cs.GT] 15 Sep 2016

Mechanism Design with Exchangeable Allocations

Qiang Zhang
csqzhang@gmail.com

Abstract

We investigate mechanism design without payments when agents have different types of preferences. Contrary to most settings in the literature where agents have the same preference, e.g. in the facility location games all agents would like to stay close to (or away from) the facility, we demonstrate the limitation of mechanism design without payments when agents have different preferences by introducing *exchanging phases*. We consider two types of exchanging phases. The first model is called *central exchanges* where the exchanges are performed by a central authority. The other model is called *individual exchanges* where agents exchange their outcomes by themselves. By using facility location games as an example, we provide a truthful mechanism that optimizes social welfare in central exchanges. We also provide a universally truthful randomized mechanism that achieves at least a half of the optimal social welfare in individual exchanges.

1 Introduction

Mechanism design aims to achieve desirable social outcomes when the partial but necessary information are possessed by selfish agents. To resolve this problem, *truthful* mechanisms are designed to provide selfish agents incentives for declaring their true private information. Incentives are usually injected by the payments in those mechanisms such as auctions. Recently, *mechanism design without payments* was proposed in [Procaccia and Tennenholtz, 2009]. They showed that, when the domain of agents' preferences is highly structured, it is possible to design truthful mechanisms without resorting to payments and achieve (approximately) good outcomes at the same time. Since then, the framework of mechanism design without payments has been successfully applied to different applications, such as voting, facility locations and machine scheduling. In this paper, we show a limitation of the framework of mechanism design without payments in some settings, especially where agents have different preferences. As there is no payment, after social outcomes (or allocations for agents) being chosen by mechanisms, selfish agents could *exchange their allocations with others to their mutual benefits*. As mechanism designers, we should care about the final social outcomes, resulting from these *exchanging phases*. In this paper, we call this problem *mechanism design with exchangeable allocations*. Now let us consider an illustrative example on facility locations games.

Facility location games. In a facility location game, the location of a facility is to be determined based on the preferences reported by selfish agents. It is often to assume that agents have certain preference structures. For example, agents have single-peaked preferences if they prefer to stay close to the facility [Procaccia and Tennenholtz, 2009], agents have obnoxious preferences if they prefer to be far away from the facility [Cheng *et al.*, 2011], or agents have double-peaked preferences if they prefer to stay at particular distance away from the facility [Filos-Ratsikas *et al.*, 2015]. In those studies, the private information of agents are their locations, which may be easily verified by the central authority. Now we assume that the locations of agents are the public information but mechanisms do not know the preference structures of agents. We consider two different types of preferences as follows.

- 'Like' preference (or single-peaked preference): agents would like to minimize their distances to the facility. We denote this preference by L.
- 'Dislike' preference (or obnoxious preference): agents would like to maximize their distances away from the facility. We denote this preference by H.

Now let us illustrate the mechanism design with exchangeable allocations by the instance in Figure 1, where five agents are located in a line segment of a length of 8 km. The locations of agents are public information. Agents are asked to report their preferences, like (L) or dislike (H). Figure 1a shows the distances between agents and their true preferences. In this case, if we follow the approach of mechanism design without payments, then it is easy to verify that locating the facility that optimizes social welfare is a truthful mechanism. That is, there is no incentive for agents to misreport their preferences regardless the preferences reported by other agents. Thus, the facility will be located at 5km (the location of agent 3) in this instance. As we mentioned before, after the facility being built, selfish agents could trade their

houses. We consider that the classical top trading cycle (TTC) algorithm [Shapley and Scarf, 1974] is used here. Therefore, agent 4 could trade her location with agent 1, which results in that agent 4 is 5km away from the facility. Now, consider the same instance except that agent 4 misreports her preference, shown as in Figure 1b. To maximize the social welfare, the facility is built at 7km (the location of agent 4). Yet, once built, agent 4 could trade her location with agent 1 so that she is 7km away from the facility. It provides her a better outcome than in Figure 1a when she reported truthfully. Readers may notice that this situation would not happen if all agents have the same preference, L or H. However, as mechanism designers, we are interested in dealing with the case when selfish agents have different preferences or even conflicted preferences. To summarize, using facility location games, we demonstrate a problem in the approach of mechanism design without payments. The insight is that mechanism designers or social planners should consider what will happen after the outcomes or allocations are chosen by mechanisms. Here we consider that agents are able to exchange their allocations for their mutual benefit.

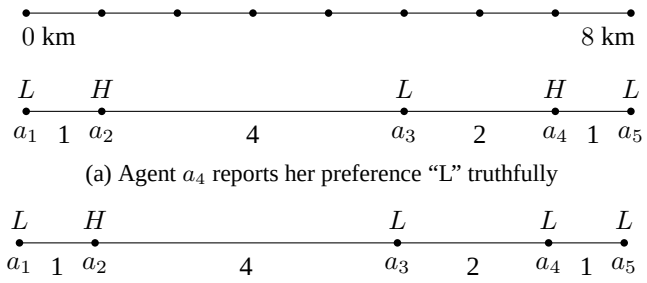

(a) Agent a_4 reports her preference "L" truthfully

(b) Agent a_4 misreports her preference as "H"

Figure 1: Example 1

Cooperative Games. Cooperative games study how agents form a coalition and benefit together. The focus of mechanism design with exchangeable allocations is *not* forming coalition. For example, in the facility location game in Figure 1, it is not necessary for agent 4 to have a consensus with agent 1 before her misreport. The assumption we made is that agents are utility maximizers. Otherwise, agent 1 may prefer to exchange her location with other agents rather than agent 4.

Top Trading Cycle. In this paper, we use the top trading cycle (TTC) algorithm [Shapley and Scarf, 1974] to simulate the exchange phase. There are different exchange protocols to trade individual good, e.g. deferred acceptance algorithm. The reasons that we choose the TTC algorithm over other choices are two-fold. First, the final outcome produced by the TTC algorithm is a core outcome, i.e., there is no other outcome in which all agents could be better off by exchanging their locations in a different way. It also implies that no exchange is possible afterwards. Second, as we are interested in truthful mechanisms, it is essential to have truthful exchange protocols. It is well-known that it is never advantageous to an agent to lie about its preference in the TTC algorithm.

Our approaches and results.

We investigate mechanism design without payments when agents are able to exchange their allocations. We first formalize the problem as mechanism design with exchangeable allocations. Then, we study it in facility location games. The first natural approach is to take exchanging phase into account when we design truthful mechanisms. In other words, we are interested in mechanisms that could exchange the locations of agents on their behalf. We call it central exchanges model. In this model, a mechanism is truthful if no agent could benefit by misreporting her preference and then exchanging her allocation with other agents. The first observation is that the design of such mechanism will highly depend on the exchange procedure. In this paper, we assume that the TTC algorithm is used in exchange phases. In addition, we consider a different benchmark. As agents may exchange their locations, better benchmarks would be the optimal social outcomes when one is able to relocate agents. In central exchanges model, we give a truthful mechanism that achieves the optimal social welfare.

Second, we understand that it is a strong assumption to allows mechanisms to relocate agents. In Section 4, we investigated the case that mechanisms only determine the location of the facility and the exchanges are carried out by agents themselves. We call it individual exchanges model. We provide a simple universally truthful randomized mechanism, which guarantees that, after agents exchanging their locations, the social welfare it achieves is at least half of the social welfare achieved by any mechanism. Note that we do not require that agents exchanges their locations according to the TTC algorithm in individual exchanges model. We only assume that agents are utility maximizers and exchange for their mutual benefit.

Related Work

Mechanism design without payments was coined in [Procaccia and Tennenholtz, 2009] to study mechanism design problems where the goal is to optimize an objective under the restriction that no money transfer is possible. The problem considered in [Procaccia and Tennenholtz, 2009] is a facility location game where selfish agents with single-peaked preferences located on a line segment would like to minimize their distances to the facility. They provided different deterministic and randomized truthful mechanisms for minimizing the total distance of agents or minimizing the maximum distance among agents. Later there is a significant amount of work on facility location games with single-peaked preferences in literature both from social choice and computer science, e.g. [Schummer and Vohra, 2002; Lu *et al.*, 2009; Lu *et al.*, 2010; Fotakis and Tzamos, 2010; Fotakis and Tzamos, 2014; Serafino and Ventre, 2015].

Besides single-peaked preferences, facility locations games are also investigated under other preferences. For example, [Filos-Ratsikas *et al.*, 2015] studied the case where agents have double-peaked preferences, in which agents would like to stay neither too close or too far away from the facility. [Zou and Li, 2015] studied facility location games

with dual preference where agents are either willing to stay close to or willing to stay away from the facility.

To the best of our knowledge, all existing work on facility location games deal with the case that the locations of agents are private information, i.e., the true locations of agents are only known to themselves. The primary reason is that most papers assume the preferences of agents and their preferences is unique. Arguably, the locations of agents are relatively easy to verify by central authority. Thus, in this paper we assume that there are two types of preference. However, the exact types of preferences are agents' private information. It is the first paper to study mechanism design on facility location games with different preferences when preferences are unknown to mechanism.

2 Mechanism Design with Exchangeable Allocations

In this section, we present the general description on mechanism design with exchangeable allocations, which consists of two phases. The first phase is similar to mechanism design without payments in [Procaccia and Tennenholtz, 2009]. There is a central authority who aims to optimize an public known objective $\mathcal{F}$ by choosing an outcome from a set of all possible outcomes denoted by O. To achieve good outcomes, the authority implements a mechanism, say $\mathcal{M}$, which first asks a set of selfish agents, denoted by A, to declare their private information and then chooses an outcome. In addition, each agent $i \in A$ has an utility function $u_i(\cdot)$ for all possible outcomes. Let c_i denote the private information of agent i and d_i be the private information declared by the agent i. Given $\mathbf{d} =< d_1, \ldots, d_n >$, Mechanism $\mathcal{M}$ chooses an outcome denoted by $\mathcal{M}(\mathbf{d})$. It is possible that agents would receive different outcomes, for example, agents would receive different resource in resource allocation problems. To allow the generality, we use $\mathcal{M}_i(\mathbf{d})$ to denote the outcome for agent i. Hence, the utility of agent i is $u_i(\mathcal{M}_i(\mathbf{d}))$.[1]

The second phase involves an (possibly pre-defined) exchange procedure $\mathcal{E}$. Agents improve their utilities by exchanging their outcomes via procedure $\mathcal{E}$. For example, agent i may exchange her outcome with agent j if both agent i and j improve their utility. That is,

$$u_i(\mathcal{M}_j(\mathbf{d})) > u_i(\mathcal{M}_i(\mathbf{d})) \text{ and } u_j(\mathcal{M}_i(\mathbf{d})) > u_j(\mathcal{M}_j(\mathbf{d}))$$

Note that the exchange procedure is not fixed and can be specified due to different purposes. For example, as we mentioned, in this paper, we are mainly interested in the top trading circle as the exchange procedure. Let $\mathcal{E}_i(\mathcal{M}(\mathbf{d}))$ be the outcome of agent i after the exchanges. Hence, in the end agent i has an utility $u_i(\mathcal{E}_i(\mathcal{M}(\mathbf{d})))$.

The goal of mechanism design with exchangeable allocations is to design mechanism that give agents incentive to declare their true private information regardless the information declared by others. Specifically , a mechanism $\mathcal{M}$ is truthful if it holds that

$$u_i(\mathcal{E}_i(\mathcal{M}(\mathbf{d}_{-i}, c_i)) \geq u_i(\mathcal{E}_i(\mathcal{M}(\mathbf{d}_{-i}, d_i))$$

for any $i, \mathbf{d}_{-i}$ and c_i where $\mathbf{d}_{-i} =< d_1, \ldots, d_{i-1}, d_{i+1}, \ldots, d_n >$.

In the rest of paper, we use facility location games to demonstrate two different approaches to design truthful mechanism.

3 Central Exchanges

In this section, we will focus on the central exchanges approach in facility location games, where mechanisms could assign new locations to agents after the location of facility has been decided. As we are interested in mechanisms that maximize social welfare, it implies that agents will have no incentives (or will not want) to exchange location themselves given their new locations.

Model

In a facility location game, a facility is to built in a line segment $[0, d]$, where n agents are located. Agent i is at point x_i that is known to the public, and has a private preference (aka, type), Like (L) or Dislike (H), denoted by t_i. We will call the collection $\mathbf{x} =< x_1, \ldots, x_n >$ a *location profile*, and the collection $\mathbf{t} =< t_1, \ldots, t_n >$ a *preference profile*. Furthermore, we say $\hat{\mathbf{x}}$ is *permuted* to x if there exists a permutation π such that for any $i \in [1, n]$ it holds $\hat{x}_i = x_{\pi(i)}$.

A deterministic mechanism is a function $\mathcal{M} : \mathbb{R}^{2n} \mapsto \mathbb{R}^{n+1}$ mapping a location profile $\mathbf{x}$ and a preference profile $\mathbf{t}$ to a permuted location profile and a point in $\mathbb{R}$ which is the location of the facility. For the purpose of easy presentation, we sometimes denote the location of the facility by $\mathcal{M}^f(\mathbf{x}, \mathbf{t})$ and the permuted location profile by $\mathcal{M}(\mathbf{x}, \mathbf{t})$. We also denote the location assigned to agent i by $\mathcal{M}_i(\mathbf{x}, \mathbf{t})$. Intuitively speaking, a mechanism builds the facility at point $\mathcal{M}^f(\mathbf{x}, \mathbf{t})$ and relocates agent i to location $\mathcal{M}_i(\mathbf{x}, \mathbf{t})$. Given the location $\hat{x}_i$ assigned to agent i and the location $y \in \mathbb{R}$ of the facility, the utility of agent i is

$$\mathrm{u}(y, \hat{x}_i, t_i) = \begin{cases} d - |\hat{x}_i - y| & \text{if } t_i = L \\ |\hat{x}_i - y| & \text{if } t_i = H \end{cases}$$

Given a permuted location profile $\hat{\mathbf{x}}$ and a location of the facility $y \in \mathbb{R}$, the social welfare is

$$\mathrm{SW}(y, \hat{\mathbf{x}}, \mathbf{t}) = \sum_{i \in [1,n]} \mathrm{u}(y, \hat{x}_i, t_i)$$

A deterministic mechanism is *truthful* if no agent would benefit by misreporting her preference regardless the locations and preferences declared by other agents. That is, for every possible location profile $\mathbf{x}$, every possible preference profile $\mathbf{t}$, every agent $i \in [1, n]$ and every $t'_i \in \{L, H\}$, it holds that $u(\mathcal{M}^f(\mathbf{x}, \mathbf{t}), \mathcal{M}_i(\mathbf{x}, \mathbf{t}), t_i) \geq u(\mathcal{M}^f(\mathbf{x}, <\mathbf{t}_{-i}, t'_i>), \mathcal{M}_i(\mathbf{x}, <\mathbf{t}_{-i}, t'_i>), t'_i)$, where $\mathbf{t}_{-i} =< t_1, \ldots, t_{i-1}, t_{i+1}, \ldots, t_n >$.

We are interested in truthful mechanisms that perform well with respect to the optimal social welfare. The performance ratio r of a mechanism is given by comparing the social welfare it achieves with the optimal social welfare on any location profile $\mathbf{x}$ and any preference profile $\mathbf{t}$ as follows.

$$r = \max_{y, \mathbf{x}, \hat{\mathbf{x}}, \mathbf{t}} \frac{\mathrm{SW}(y, \hat{\mathbf{x}}, \mathbf{t})}{\mathrm{SW}(\mathcal{M}^f(\mathbf{x}, \mathbf{t}), \mathcal{M}(\mathbf{x}, \mathbf{t}), \mathbf{t})}$$

[1] We assume there is no externalities.

where $\hat{\mathbf{x}}$ is permuted to $\mathbf{x}$.

Optimal Mechanism

We first present a customized location permutation (see Algorithm 1), which returns a location profile that is permuted to $\hat{\mathbf{x}}$ given location profiles $\mathbf{x}$ and $\hat{\mathbf{x}}$, a preference profile $\mathbf{t}$ and a facility location y.

Algorithm 1: LocationPermuation($\mathbf{x}, \hat{\mathbf{x}}, \mathbf{t}, y$)

1 Let $\mathbf{L}$ and $\mathbf{H}$ denote the set of agents whose preferences are L and H, respectively;
2 Let $\hat{\mathbf{x}}_{\mathbf{L}}$ and $\hat{\mathbf{x}}_{\mathbf{H}}$ denote the locations of agents in $\mathbf{L}$ and $\mathbf{H}$ in $\hat{\mathbf{x}}$, respectively;
3 Sort agents in L by their distances to y in $\mathbf{x}$ in an increasing order. Similarly, sort agents in H by their distances to y in $\mathbf{x}$ in a decreasing order;
4 Sort locations in $\hat{\mathbf{x}}_{\mathbf{L}}$ by the distances to y in an increasing order. Similarly, sort locations in $\hat{\mathbf{x}}_{\mathbf{H}}$ by the distances to y in a decreasing order;
5 Assign agents in L sequentially to locations in $\hat{\mathbf{x}}_{\mathbf{L}}$ based on their orders, and similarly assign agents in H sequentially to locations in $\hat{\mathbf{x}}_{\mathbf{H}}$;
6 Return the assignment.

Essentially, Algorithm 1 assigns the locations $\hat{\mathbf{x}}_{\mathbf{L}}$(or $\hat{\mathbf{x}}_{\mathbf{H}}$) to agents in $\mathbf{L}$(or $\mathbf{H}$) according to their distances to the facility in location profile $\mathbf{x}$. Now, we present the following mechanism:

Mechanism $\mathcal{OPT}$. *Given a location profile $\mathbf{x}$ and a preference profile $\mathbf{t}$, choose a permuted locations $\mathbf{z}$ and y such that $\mathrm{SW}(y, \mathbf{z}, \mathbf{t})$ is maximized, tie-breaking arbitrarily. The mechanism returns y as the location of the facility $\mathcal{OPT}^f(\mathbf{x}, \mathbf{t})$, and call $\mathrm{LocationPermuation}(\mathbf{x}, \mathbf{z}, \mathbf{t}, \mathrm{y})$ to compute the locations for agents $\mathcal{OPT}(\mathbf{x}, \mathbf{t})$.*

Note that z and y may not be unique. We assume that Mechanism $\mathcal{OPT}$ uses an arbitrary tie-breaking rule. The tie-breaking rule has no effect on the performance of the mechanism as it always maximizes the social welfare. As we will show that, if agents could benefit by lying their preferences in Mechanism $\mathcal{OPT}$, then the social welfare is not maximized. Therefore, any arbitrary tie-breaking rule works for the truthfulness.

Examples of Algorithm 1. Now we would like to illustrate Mechanism $\mathcal{OPT}$ using the instance in Figure 1. Readers could verify that $\mathrm{SW}(y, \mathbf{z}, \mathbf{t})$ is maximized when $y = 8$ and $\mathbf{z} = \{z_1 = 5, z_2 = 1, z_3 = 8, z_4 = 0, z_5 = 7\}$. Mechanism $\mathcal{OPT}$ locates the facility at 8. Now let us see the locations assigned to agents. As agent 5 is close to the location of facility among agents with preference L, Mechanism $\mathcal{OPT}$ assigns agent 5 to location 8. As agent 3 is the second closest to the location of facility after agent 5, Mechanism $\mathcal{OPT}$ assigns agent 3 to location 7. Next, Mechanism $\mathcal{OPT}$ assigns agent 1 to location 5. Similarly, for agents with preference H, Mechanism $\mathcal{OPT}$ assigns agent 4 to location 1 and agent 2 to location 0.

Computing $\mathbf{z}$ in $\mathcal{OPT}$. Mechanism $\mathcal{OPT}$ first chooses the location of the facility y and a location profile $\mathbf{z}$ that is permuted to $\mathbf{x}$ such that the social welfare is maximized. In the following we show that it is computationally feasible to find a y and $\mathbf{z}$ given $\mathbf{x}$ and $\mathbf{t}$.

Lemma 1. *Given $\mathbf{x}$, $\mathbf{t}$, y and $\mathbf{z}$ that is permuted to $\mathbf{x}$, if $\mathrm{SW}(y, \mathbf{z}, \mathbf{t})$ is maximized, then for any $i \in \mathbf{H}$, $z_i \leq \min \mathbf{z}_{\mathbf{L}}$ or $z_i \geq \max \mathbf{z}_{\mathbf{L}}$.*

Intuitively, Lemma 1 implies that agents whose preferences are L are assigned to consecutive locations where social welfare is maximized. This characterization strictly reduces the space of the optimal outcomes. Next we present another characterization on the locations of the facility in the optimal outcomes.

Lemma 2. *Given $\mathbf{x}$, $\mathbf{t}$ and $\mathbf{z}$ that is permuted to $\mathbf{x}$, there exists a y such that $y \in \mathbf{x} \cup \{0, d\}$ and $\mathrm{SW}(y, \mathbf{z}, \mathbf{t})$ is maximized.*

By Lemma 1 and 2, we limit the search space for the optimal outcomes to $O(n^2)$. The optimal outcomes can be found by enumerating all possible outcomes and returning the outcome with the maximum social welfare. If ties exist, Mechanism $\mathcal{OPT}$ breaks ties arbitrarily.

Lemma 3. *Mechanism $\mathcal{OPT}$ is truthful.*

Proof. We show that, given any $\mathbf{x}, \mathbf{t}$, any agent $i \in H$ cannot increase her utility by declaring $t_i = L$. The proof for the opposite case that any agent $i \in L$ cannot increase her utility by declaring $t_i = H$ is symmetric and omitted due to the space limit.

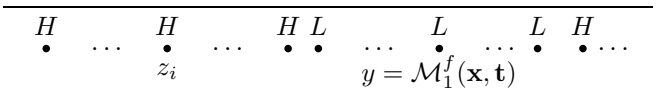

(a) The outcome when agent a_i reports her preference "H" truthfully in $\mathcal{M}_1$

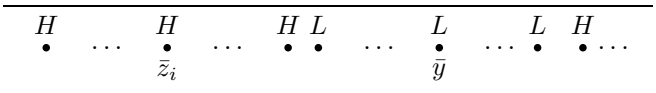

(b) A possible outcome when agent a_i reports her preference as L

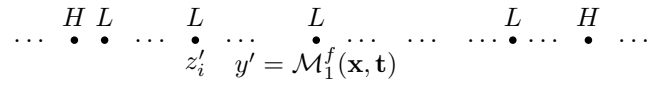

(c) The outcome the maximize social welfare when agent a_i misreports her preference as L in $\mathcal{M}_1$

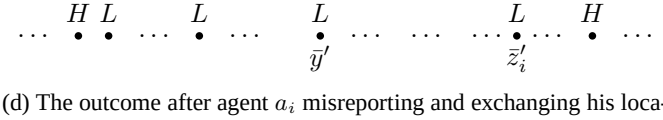

(d) The outcome after agent a_i misreporting and exchanging his location with other agent

Figure 2: Instances in the proof of Lemma 3

Before presenting the proof, we introduce the following four instances that will be used in this proof. The instances are shown in Figure 2,

- Instance a: $\mathbf{x}$ and $\mathbf{t}$. Outcome: $< y, \mathbf{z} >$. In this instance, agents report preference truthfully. $\mathcal{M}_1$ places the facility at y and reassigns each agent $j \in A$ to z_j. The social welfare is maximized.
- Instance b: $\bar{\mathbf{x}}$ and $\bar{\mathbf{t}}$ where $\bar{\mathbf{x}} = \mathbf{x}$ and $\bar{\mathbf{t}} = \mathbf{t}$ except $\bar{t}_i = L$. Outcome: $< \bar{y}, \bar{\mathbf{z}} >$ where $\bar{y} = y$ and $\bar{\mathbf{z}} = \mathbf{z}$. In this instance, agent i misreports her preference. The facility is placed at $\bar{y}$ and agent $j \in A$ is assigned to $\bar{z}_j$.
- Instance c: $\mathbf{x}'$ and $\mathbf{t}'$ where $\mathbf{x}' = \bar{x}$ and $\mathbf{t}' = \bar{\mathbf{t}}$. Outcome: $< y', \mathbf{z}' >$. In this instance, Mechanism $\mathcal{OPT}$ places the facility at y' and assigns each agent $j \in A$ to z'_j. The social welfare is maximized.
- Instance d: $\bar{\mathbf{x}}'$ and $\bar{\mathbf{t}}'$ where $\bar{\mathbf{x}}' = x$ and $\bar{\mathbf{t}}' = \mathbf{t}$. Outcome: $< \bar{y}', \bar{\mathbf{z}}' >$. This instance is essentially the exchange phase, where agent i exchanges his location in $\mathbf{z}'$ with agent j who is furthest to y' with preference L in $\mathbf{z}'$.

From these instances , we get the following:

$$\begin{aligned}
&\sum_{j\in H\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) + u(\bar{y},\bar{z}_i,L) \\
\le &\sum_{j\in H\setminus\{i\}} u(y',z'_j,t'_i) + \sum_{j\in L\setminus\{i\}} u(y',z'_j,t'_i) + u(y',z'_i,L) \\
= &\sum_{j\in H\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i) + u(\bar{y}',\bar{z}'_i,L)
\end{aligned} \tag{1}$$

The first inequality comes from that y' and $\mathbf{z}'$ maximizes the social welfare when there are $|L|+1$ agents with preference L. [2] Now let us assume that agent i could benefit by misreporting her preference as L and then exchanging her location with other agent. As $u(\bar{y},\bar{z}_i,H) = u(y,z_i,H)$ by definition, it implies that $u(\bar{y}',\bar{z}'_i,H) > u(\bar{y},\bar{z}_i,H)$. Since the relation of two types of preferences are opposite, we know that that $u(\bar{y}',\bar{z}'_i,L) < u(\bar{y},\bar{z}_i,L)$. Substitute it into Equation (1), we get

$$\begin{aligned}
&\sum_{j\in H\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) \\
< &\sum_{j\in H\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i)
\end{aligned} \tag{2}$$

Hence, we have

$$\begin{aligned}
&\sum_{j\in H\setminus\{i\}} u(y,z_j,t_i) + \sum_{j\in L\setminus\{i\}} u(y,z_j,t_i) + u(y,z_i,H) \\
= &\sum_{j\in H\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y},\bar{z}_j,\bar{t}_i) + u(\bar{y},\bar{z}_i,H) \\
< &\sum_{j\in H\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i) + \sum_{j\in L\setminus\{i\}} u(\bar{y}',\bar{z}'_j,\bar{t}'_i) + u(\bar{y}',\bar{z}'_i,H)
\end{aligned}$$

It contradicts the fact that Mechanism $\mathcal{OPT}$ maximizes the social welfare. □

[2] In instances b and c the preference of agent i is L

4 Individual Exchanges

In this section, we explore the truthful mechanisms when social planners are less powerful. In particular, exchanges between agents are coordinated by agents themselves rather than mechanisms. We assume that the TTC algorithm simulates the exchanges between agents. The problem is formalized as follows.

Model

The goal is to locate a facility in a line segment $[0, d]$ given a location profile $\mathbf{x}$ and a preference profile $\mathbf{t}$, which are defined as previous section.

A deterministic mechanism is a function $\mathcal{M} : \mathbb{R}^{2n} \to \mathbf{R}$ which is the location of the facility. A randomized mechanism is a function $\mathcal{M} : \mathbb{R}^{2n} \to \Delta(\mathbf{R})$ where $\Delta(\mathbf{R})$ is the set of distribution over $\mathbf{R}$. Given any location profile $\mathbf{x}$ and preference profile $\mathbf{t}$, let $\mathcal{M}(\mathbf{x},\mathbf{t})$ denote the location of facility in mechanism $\mathcal{M}$. $\mathcal{M}(\mathbf{x},\mathbf{t})$ is a random variable in a randomized mechanism.

Given the location of the facility y, the location profile $\mathbf{x}$ and the preference profile $\mathbf{t}$, let $\mathcal{TTC}(\mathbf{x},\mathbf{t},y)$ denote the outcome of the TTC algorithm. In particular, let $\mathcal{TTC}_i(\mathbf{x},\mathbf{t},y)$ be the location of agent i from the TTC algorithm. Hence, the utility of agent i is

$$\mathrm{u}(y,\mathcal{TTC}_i(\mathbf{x},\mathbf{t},y),t_i) = \begin{cases} d - |\mathcal{TTC}_i(\mathbf{x},\mathbf{t},y) - y| & \text{if } t_i = L \\ |\mathcal{TTC}_i(\mathbf{x},\mathbf{t},y) - y| & \text{if } t_i = H \end{cases}$$

Similarly, the social welfare obtained by TTC is

$$\mathrm{SW}(y,\mathcal{TTC}(\mathbf{x},\mathbf{t},y),\mathbf{t}) = \sum_{i\in[1,n]} \mathrm{u}(y,\mathcal{TTC}_i(\mathbf{x},\mathbf{t},y),t_i)$$

A deterministic mechanism $\mathcal{M}$ is *truthful* if no agent would benefit by misreporting her preference in the mechanism and then obtaining a better outcome from the TTC algorithm regardless of the locations and preferences of other agents. That is, for any location profile $\mathbf{x}$, any preference profile $\mathbf{t}$, any $i \in [1,n]$ and any $t'_i \in \{L,H\}$, it holds $u(\mathcal{M}(\mathbf{x},\mathbf{t}),\mathcal{TTC}_i(\mathbf{x},\mathbf{t},\mathcal{M}(\mathbf{x},\mathbf{t})),t_i) \ge u(\mathcal{M}(\mathbf{x},<\mathbf{t}_{-i},t'_i>),\mathcal{TTC}_i(\mathbf{x},<\mathbf{t}_{-i},t_i>,\mathcal{M}(\mathbf{x},<\mathbf{t}_{-i},t'_i>)),t_i)$. Note that the preference t'_i misreported by agent i only affects the location of the facility chosen by mechanism $\mathcal{M}$ only. Since the exchanges happened in the TTC algorithm are carried out by agents themselves, the true preference t_i is used there. A universally truthful mechanism is a probability distribution over deterministic truthful mechanisms

As the previous section, we measure the performance of the truthful mechanism with its approximation ratio to the optimal social welfare. Given mechanism $\mathcal{M}$, its approximation ratio is defined as follows.

$$r^{\mathcal{M}} = \max_{y,\mathbf{x},\hat{\mathbf{x}},\mathbf{t}} \frac{\mathrm{SW}(y,\hat{\mathbf{x}},\mathbf{t})}{\mathrm{SW}(\mathcal{M}(\mathbf{x},\mathbf{t}),\mathcal{TTC}(\mathbf{x},\mathcal{M}(\mathbf{x},\mathbf{t}),\mathbf{t}),\mathbf{t})}$$

where $\hat{\mathbf{x}}$ is permuted to $\mathbf{x}$.

Optimal location

From the previous section, we know that, given the location of the facility, agents with the same preference (L or H) will

be consecutive after their exchanges in the TTC algorithm. Hence, it is easy to verify that, assuming agents report their preference truthfully, one could locate the facility at the optimal location y so that the optimal social welfare will be achieved after the exchanges between agents. Hence, the following question arises naturally: *Is it a truthful mechanism by locating the facility at the optimal location conditioned on that agents would exchange their locations by TTC algorithm?* Note that in the example provided in Figure 1 the facility is located at the optimal location assuming that agents *would not* exchange their locations.

Theorem 1. *Given* $\mathbf{x}$ *and* $\mathbf{t}$, $\mathcal{M}(\mathbf{x},\mathbf{t}) = \arg\max_{y\in[0,d]} \mathrm{SW}(y,\hat{\mathbf{x}},\mathbf{t})$ *is not truthful given that agents would exchange their locations by the TTC algorithm.*

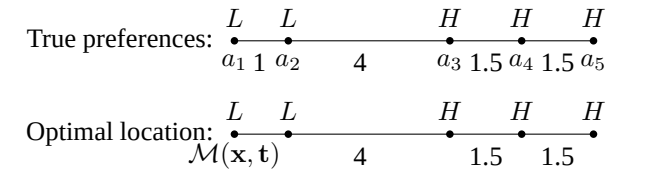

(a) Agent a_4 reports her preference "H" truthfully

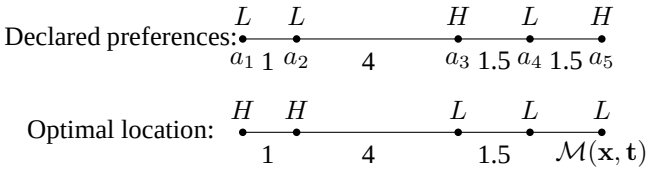

(b) Agent a_4 misreports her preference as "L"

Figure 3: Instances in Theorem 1

Proof. We prove this theorem by the instances shown in Figure 3 in which five agents are located in a line segment of a length of 8km. Figure 3a shows the distance between agents and their true preferences. The optimal location of the facility is at 0km (the location of agent 1). No exchange would happen in the TTC algorithm as no pair of agents would prefer exchanging the locations. The utility of agent 4 is 6.5. Now considering the case that agent 4 falsely declares her preference as L shown in Figure 3b, the optimal location of the facility is at 8km. Then in the TTC algorithm, agent 5 would exchange her location with agent 1, and agent 4 would exchange her location with agent 2. The utility of agent 4 is 7. This implies that agent 4 benefit by manipulating her preference in the mechanism. Then the theorem directly follows. □

Simple 2-approximation Mechanism

Next we provide a simple universally truthful randomized mechanism that achieve at least half of the optimal social welfare in expectation.

Mechanism $\mathcal{M}$. *Given a location profile* $\mathbf{x}$ *and a preference profile* $\mathbf{t}$, $\mathcal{M}^f(\mathbf{x},\mathbf{t}) = 0$ *with probability* $1/2$*, and* $\mathcal{M}^f(\mathbf{x},\mathbf{t}) = d$ *with probability* $1/2$*.*

Theorem 2. *Mechanism $\mathcal{M}$ is a universally truthful randomized mechanism that achieves* 2*-approximation to the optimal social welfare.*

Universally truthfulness

It is easy to verify that agents would not benefit from misreporting their preference as Mechanism $\mathcal{M}$ always places the facility at location 0 with a probability of $1/2$ and places the facility at location d with a probability of $1/2$.

Approximation

The proof of the approximation of Mechanism $\mathcal{M}$ comes from two ingredients.

Lemma 4. *During exchanging phases, the social welfare increases as agents exchange their locations with each other.*

This lemma allows us to only consider the social welfare before agents exchanging their locations. In other word, if the social welfare before exchanging phases is close to the optimal social welfare, then the social welfare after exchanging phases is still close to the optimal social welfare. With some simple calculations, we get the following.

Lemma 5. *Given a location profile* $\mathbf{x}$ *and a preference profile* $\mathbf{t}$, $\mathrm{SW}(\mathcal{M}^f(\mathbf{x},\mathbf{t}),\mathbf{x},\mathbf{t}) = \frac{n\cdot d}{2}$.

Proof. As Mechanism $\mathcal{M}$ is randomize mechanism, the expected social welfare it achieves can be written as follows.

$$\begin{aligned}
&\mathrm{SW}(\mathcal{M}^f(\mathbf{x},\mathbf{t}),\mathbf{x},\mathbf{t})\\
=&\mathrm{SW}(\mathbb{E}[\mathcal{M}^f(\mathbf{x},\mathbf{t})],\mathbf{x},\mathbf{t})\\
=&\frac{1}{2}\mathrm{SW}(0,\mathbf{x},\mathbf{t}) + \frac{1}{2}\mathrm{SW}(d,\mathbf{x},\mathbf{t})\\
=&\frac{1}{2}\sum_{i\in A} u(0,x_i,t_i) + \frac{1}{2}\sum_{i\in A} u(d,x_i,t_i)
\end{aligned}$$

By the fact that $u(0,x_i,t_i) + u(d,x_i,t_i) = d$ for any agent in A regardless x_i and t_i, it concludes that $\mathrm{SW}(\mathcal{M}^f(\mathbf{x},\mathbf{t}),\mathbf{x},\mathbf{t}) = \frac{n\cdot d}{2}$. □

By the fact that $\mathrm{SW}(\mathcal{M}(\mathbf{x},\mathbf{t}),\mathcal{TTC}(\mathbf{x},\mathcal{M}(\mathbf{x},\mathbf{t}),\mathbf{t}),\mathbf{t}) > \mathrm{SW}(\mathcal{M}(\mathbf{x},\mathbf{t}),\mathbf{x},\mathbf{t})$ and the utility of each agent is at most d, the approximation of Mechanism $\mathcal{M}$ directly follows.

5 Conclusion

In this paper, we consider the mechanism design without payments when selfish agents have difference types of preferences. We identify a limitation of the classical approach, that is, agents could benefit by exchanging their allocations obtained from mechanisms. To cope with this issue, we present a model called mechanism design with exchangeable allocations. We consider two approaches for this model, central exchanges and individual exchanges. By using facility location games as an example, we provide a truthful mechanism that achieves the optimal social welfare in central exchanges. We also provide a universally truthful randomized mechanism that achieves at least half of the optimal social welfare in individual exchanges. It will be interesting to investigate the problem in settings other than facility location games. We leave it as future work.

References

[Cheng *et al.*, 2011] Yukun Cheng, Wei Yu, and Guochuan Zhang. Mechanisms for obnoxious facility game on a path. In *Combinatorial Optimization and Applications*, pages 262–271. Springer, 2011.

[Filos-Ratsikas *et al.*, 2015] Aris Filos-Ratsikas, Li Minming, Jie Zhang, and Qiang Zhang. Facility location with double-peaked preference. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, 2015.

[Fotakis and Tzamos, 2010] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, pages 234–245, 2010.

[Fotakis and Tzamos, 2014] Dimitris Fotakis and Christos Tzamos. On the power of deterministic mechanisms for facility location games. *ACM Transactions on Economics and Computation*, 2(4):15, 2014.

[Lu *et al.*, 2009] P. Lu, Y. Wang, and Y. Zhou. Tighter bounds for facility games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, pages 137–148, 2009.

[Lu *et al.*, 2010] P. Lu, X. Sun, Y. Wang, and Z.A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM conference on Electronic commerce (ACM-EC)*, pages 315–324, 2010.

[Procaccia and Tennenholtz, 2009] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on Electronic commerce (ACM-EC)*, pages 177–186. ACM, 2009.

[Schummer and Vohra, 2002] J. Schummer and R.V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2002.

[Serafino and Ventre, 2015] Paolo Serafino and Carmine Ventre. Truthful mechanisms without money for non-utilitarian heterogeneous facility location. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, 2015.

[Shapley and Scarf, 1974] Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.

[Zou and Li, 2015] Shaokun Zou and Minming Li. Facility location games with dual preference. In *Proceedings of 14th Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 615–623, 2015.